Real-world Network Automation

Matt Peterson – Cumulus Networks

February 4th, 2015 – NANOG63

Network automation landscape

Panelists

- Intros
- Statements

Q&A

- Prompted
- Audience

Traditional vendor options

- CLI screen scraping
- NETCONF/YANG (only multi-vendor option)
- XML via CLI
- REST'ful API

Upstarts (all SW vendors)

- Cumulus – "Linux unencumbered" (treat as a server)
- Pica8 – "Choose your own adventure" (L3, OpenFlow)
- Bigswitch ONL – "OpenWRT for bare metal" (limited fwd'ing)

What is driving all this?

- DevOps as a business agility enabler
- Highly dynamic environments (VMs spin up, down, migration)
- Rancid (usually one-way sync) – no longer acceptable

What has changed?

- Users demand for multi-vendor, simple (sorry NETCONF)
- Compute tools as native agents under network OSs
- Tightly coupled workflow, network unacceptable as a silo

Carlos Vicente

Network Automation Panel
NANOG 63

About Carlos

- Network Engineer at Dyn
  - Network automation one of his main priorities
- Previously ISC, NSRC, Univ. of Oregon
- Not afraid to write code
- Likes Open Source
  - Author of Netdot (netdot.uoregon.edu)
- From the Dominican Republic
  - Lives in New Hampshire now

About Dyn

- Internet Performance
  - Advanced DNS services
  - E-mail delivery
  - Internet intelligence
- Started in New Hampshire
  - Becoming a global company
- ~20 data centers in 5 continents
  - Hundreds of network devices

Automation at Dyn

- Existing compute, application automation
  - Continuous Integration, etc.
  - Growing and improving fast in this area
- As of last year, no automation on network side
  - Project Kipper in early stages, but already in production

Provisioning Automation

Bronwyn Lewis
NANOG63
San Antonio, Texas

Bronwyn Lewis

- Worked in operations, engineering, development, and as a technical writer & project manager in tech and entertainment research (~7 years)
- Studied international affairs, human rights issues, and governance at The New School in NYC (~3 years)
- Provisioning engineer at Packet Clearing House since November 2013 (~15 months)
- First Ansible playbook: August 2014


```
➜  provisioning git:(master) ansible-playbook provisioning.yml

PLAY [Generate configuration files] ************************

TASK: [router | Generate configuration files] **************
ok: [localhost] => (item={'as3856': None, 'as42_v6': '2806:239
```

Packet Clearing House

- 501(c)3 non-profit based in the Bay Area, known for supporting operations and analysis in the areas of internet traffic and routing, as well as supporting IXPs
- Hosts multiple root nameserver mirrors
- DNS anycast for ~150 ccTLDs & gTLDs
- ~100 locations globally (including Sudan & Vanuatu)
- ~90% Cisco equipment, on the way to being 100%
- Upgraded or newly deployed to ~40 PoPs in 2014

Current + Future Automation

Current

- Ansible templating for new sites (~25) & most common equipment (4 models)
- PXE/Kickstart for server provisioning
- Python for BIOS/CIMC firmware upgrades on servers (work in progress)

Future

- All sites (>100) & equipment (~14 models) templated
- Further automate server provisioning using Ansible
- Provisioning and NOC automation tools integrated (Ansible? Schprokits?)

CLOUDFLARE®

Network Automation

Jérôme Fleury

NANOG63

04.02.2015

Who am I ?

Network engineer with some large scale automation experience: Python, Netconf, REST APIs for network provisioning

- 1500 routers in 3 years for Local Loop in France (2003-2006)
  - first experience with automation: Perl script reading Excel file from project managers, generating templates based config

Had to deploy some automation by necessity:

- you can't deploy 1500 routers in 3 years without automating the generation of the configurations

Who are we ?

Fast growing CDN/Security company:

- 32 POPs worldwide and counting
- 2000+ eBGP peering sessions
- Hundreds of Flowspec rules added every week
- thousands of servers all configured to do the same job: serving HTTP(s) and DNS requests
- But routers are all different: different vendors, different performances, different routing policies

Past and present of automation

- Frameworks (Django, Ruby on Rails) make it easier to integrate an ORM and web based views
- Devops best practices are reaching Netops: end of the lonesome Perl coder
- Traditional vendors offer vary degrees of Netconf support, while start-ups/ SDN companies are leaning towards a Devops approach: REST APIs, JSON, instead of the bulky RPC/XML
- Companies still need to develop their own tools: there's a lack of a common , vendor-independent, open-source API

Q&A - Cultural

Getting peers (people, not ASN's) to adopt tools?

- Automation has high level of up-front investment (ie: choose framework, write templates, documentation, training, etc.)
- Prerequisites: Linux, git, YAML, scripting

How does this extend to NOC (break/fix), BizOps (billing) – groups outside of operational responsibility?

Q&A - Tools

{Compute} DevOps tools assume DevOps workflow

- Version control system / central “source of truth”
- Build / unit testing / QA verification
- Monitoring

Does this approach fit with an service provider world?

Multi-vendor is difficult

- NETCONF not a 1st class citizen
- XML considered “old school” or legacy
  Some implementations receive little vendor QA treatment
- What alternatives – adapt to status quo, define a new standard and/or schema, …?

Ideal and/or realistic vision of the future?

- Schprokits – “Ansible tuned for NetOps”
- Cumulus/Pica8/ONL – “bare metal Linux as compute”
- Multi-vendor REST schema or some standards effort

Questions